arXiv:1703.09772v1 [stat.ML] 28 Mar 2017

Particle Filtering for PLCA model with Application to Music Transcription

Cazau Dorian[a] [1], Guillaume Revillon[a], Yuancheng Wang[a] and Olivier Adam[a]

[a] Sorbonne Universités, UPMC University Paris 06/CNRS, UMR 7190, Institut Jean le Rond, d'Alembert, F-75015, Paris, France

[1]Corresponding author e-mail: cazaudorian@outlook.fr

Abstract

Automatic Music Transcription (AMT) consists in automatically estimating the notes in an audio recording, through three attributes: onset time, duration and pitch. Probabilistic Latent Component Analysis (PLCA) has become very popular for this task. PLCA is a spectrogram factorization method, able to model a magnitude spectrogram as a linear combination of spectral vectors from a dictionary. Such methods use the Expectation-Maximization (EM) algorithm to estimate the parameters of the acoustic model. This algorithm presents well-known inherent defaults (local convergence, initialization dependency), making EM-based systems limited in their applications to AMT, particularly in regards to the mathematical form and number of priors. To overcome such limits, we propose in this paper to employ a different estimation framework based on Particle Filtering (PF), which consists in sampling the posterior distribution over larger parameter ranges. This framework proves to be more robust in parameter estimation, more flexible and unifying in the integration of prior knowledge in the system. Note-level transcription accuracies of 61.8 % and 59.5 % were achieved on evaluation sound datasets of two different instrument repertoires, including the classical piano (from MAPS dataset) and the marovany zither, and direct comparisons to previous PLCA-based approaches are provided. Steps for further development are also outlined.

1 Introduction

1.1 Background on PLCA

Probabilistic Latent Component Analysis (PLCA) is a straightforward extension of Probabilistic Latent Semantic Indexing (Hofmann, 1999) which deals with an arbitrary number of dimensions and can exhibit various features such as sparsity or shift-invariance. The basic model is defined as

$$P(x) = \sum_{z} P(z) \prod_{j=1}^{J} P(x_j|z) \tag{1}$$

where $P(x)$ is an J-dimensional distribution of the random variable $x = (x_1, \ldots, x_J)$, z is a latent variable and the $P(x_j|z)$ are one dimensional distributions with $j \in \{1, \ldots, J\}$. Such a general model has been successfully applied to audio signals, with a theoretical framework developed by (Smaragdis et al., 2006). Especially, PLCA has been proven to be an efficient probabilistic tool for non-negative data analysis, which offers a convenient way of designing spectrogram models. From its general formulation (eq. 1), and considering a spectrogram $S(f,t)$ as a probability distribution $P(f,t)$, a latent variable z is introduced to model $P(f,t)$ as

$$P(f,t) = \sum_{z} P(z)P(f|z)P(t|z) = \sum_{z} P(z,t)P(f|z) \tag{2}$$

where f and t represent respectively frequency and time, and are both conditionally independent given z, $P(f|z)$ are the spectral bases corresponding to component z, and $P(z,t)$ their time activations. Since there is usually no closed-form solution for the maximization of the log-likelihood or the posterior, iterative update rules based on the Expectation-Maximization (EM) algorithm are employed to estimate $P(f|z)$ and $P(z,t)$.

1.2 Limitations of current PLCA models

The major limitation of current PLCA models lies in the inherent problems of the EM algorithm. This algorithm was originally introduced by (Dempster et al., 1977) to overcome the difficulties in maximizing likelihoods of missing data models. The main advantage of that method is its easy implementation,

consisting of initializing the parameters and iterating expectation and maximization likelihoods in a step-by-step process until convergence. Its major drawback, besides the requirement of convex likelihoods, lies in its sensitivity to initialization, which increase the risks to local convergences (Robert and Casella, 1999). That issue is exacerbated in the case of multimodal likelihoods. Indeed, the increase of the likelihood function at each step of the algorithm ensures its convergence to the maximum likelihood estimator in the case of unimodal likelihoods, but implies a dependence on initial conditions for multimodal likelihoods. Alternative techniques have also been proposed to optimize the search of global maxima, such as running the algorithm a number of times with different, random starting points, or using variants from the basic EM algorithm such as Deterministic Annealing EM (DAEM) algorithm (Ueda and Nakano, 1998). These theoretical issues have reached research fields working on audio signals. To tackle the problem of dependency to initialization, some authors (Grindlay and Ellis, 2010; Benetos and Dixon, 2013) perform a training of the instrument templates, which has proved to be an effective way to initialise the spectral bases. Indeed, by fixing them without data-driven updating, we obtain a stable output for the gain function, independent of its initialisation. However, when the model becomes more complex with for example the introduction of different instrument variables, performing robust initialization is more difficult. For what concerns the local convergence problem, some works (Hoffman et al., 2009; Grindlay and Ellis, 2010; Cheng et al., 2013) have used the DAEM algorithm based on a temperature parameter.

This limitation becomes particularly critical when integrating priors into the PLCA framework. Generally speaking, this integration introduces generic problems in optimization convergence to global maxima, especially when the prior has a multi-modality form. Indeed, when a prior is injected, the maximization step becomes a maximum a posteriori step and the log posterior probability needs to have the right properties for maximization. (Fuentes et al., 2013) used of a numerical fixed point algorithm to solve the modified EM equations with a sparsity prior, whose convergence is only theoretically supposed, but "observed in practice" (although the sensitivity of the algorithm convergence to the evaluation sound dataset is not detailed). (Benetos and Dixon, 2013) privileged the use of pre-defined templates, which allows them to skip computing the EM update equation of templates, and just to apply a sparsity constraint on the pitch activity matrix and the pitch-wise source contribution matrix. Also, the simultaneous use of several priors on

a same model parameter leads to some difficulties in terms of mathematical calculation and increases convergence problems (Fuentes et al., 2013).

1.3 Particle filtering

In the framework of Bayesian variable selection, Markov Chain Monte Carlo, or Particle filtering (PF), type approaches have been proposed (Févotte and Godsill, 2006a; Févotte et al., 2008). These methods consist in sampling the posterior distribution over larger parameter ranges, making them more demanding than their EM-like counterparts, but which also, in return, offer increased robustness in convergence (i.e. reduced problems of convergence to local minima) and a complete Monte Carlo description of this parameter posterior density (Févotte and Godsill, 2006b,a; Févotte et al., 2008).

1.3.1 General Overview

Many problems in statistical signal processing (Fong et al., 2002; Andrieu et al., 2003; Vermaak et al., 2000) can be stated in a state space form as follows,

$$x_{t+1} \sim f(x_{t+1}|x_t) \tag{3}$$

$$y_{t+1} \sim g(y_{t+1}|x_{t+1}) \tag{4}$$

where $\{x_t\}$ are unobserved states of the system and $\{y_t\}$ are observations made over some time, t. $f(.|.)$ and $g(.|.)$ are pre-specified state evolution and observation densities. A primary concern in many state-space inference problems is the sequential estimation of the filtering distribution $p(x_t|y_{1:t})$, and the simulation of the entire smoothing distribution $p(x_{1:t}|y_{1:t})$, where $y_{1:t} = (y_1, y_2, \cdots, y_t)$ and $x_{1:t} = (x_1, x_2, \cdots, x_t)$. Updating of the filtering distribution can be achieved, in principle, using the standard filtering recursions (Robert and Casella, 1999)

$$p(x_{t+1}|y_{1:t}) = \int p(x_t|y_{1:t}) f(x_{t+1}|x_t) dx_t \tag{5}$$

$$p(x_{t+1}|y_{1:t+1}) = \frac{g(y_{t+1}|x_{t+1}) p(x_{t+1}|y_{1:t})}{p(y_{t+1}|y_{1:t})} \tag{6}$$

Smoothing can also be performed recursively backwards in time using the smoothing formula (Robert and Casella, 1999)

$$p(x_t|y_{1:T}) = \int p(x_{t+1}|y_{1:T}) \frac{p(x_t|y_{1:t}) f(x_{t+1}|x_t)}{p(x_{t+1}|y_{1:t})} dx_t \tag{7}$$

In practice, these filtering (eq. 5) and smoothing (eq. 7) computations can only be performed in closed form for linear Gaussian models using the Kalman filter / smoother, and for finite state-space hidden Markov models. In the case of non-linear non-Gaussian models, there is no general analytic expression for the computations of these density functions. As a consequence, an approximation strategy is required to estimate the filtering and smoothing densities, which is commonly performed with the PF method, also known as sequential Monte Carlo methods. Within the PF framework, the filtering distribution is approximated with an empirical distribution formed from point masses also called particles,

$$p(x_t|y_{1:t}) \approx \sum_{i=1}^{N} w_t^{(i)} \delta(x_t - x_t^{(i)}) \tag{8}$$

$$\sum_{i=1}^{N} w_t^{(i)} = 1, w_t^{(i)} \geq 0 \tag{9}$$

where $\delta(.)$ is the Dirac delta function and $w_t^{(i)}$ is a weight attached to particle $x_t^{(i)}$. Given this particle approximation to the posterior distribution, we can estimate the expected value of any function f $w.r.t$ the distribution $I(f,t)$, defined as $I(f_t) = \int f(x_t) p(x_t|y_{1:t}) dx_t$, using the following Monte Carlo approximation

$$I(f_t) \approx \sum_{i=1}^{N} f(x_t^{(i)}) w_t(i) \tag{10}$$

Particle smoothers generate batched realisations of $p(x_{1:T}|y_{1:T})$ based on the forward PF results. In other words, the particle smoothers are an efficient method for generating realisations from the entire smoothing density $p(x_{1:T}|y_{1:T})$ using filtering approximation.

1.3.2 Filtering

We consider the filtering distribution $p(x_t|y_{1:t})$. Using the Bayes' rule, this distribution can be rewritten as follows,

$$p(x_t|y_{1:t}) = p(x_t|y_t, y_{1:t-1}) \tag{11}$$
$$\propto p(y_t|x_t, y_{1:t-1})p(x_t|y_{1:t-1}) \tag{12}$$
$$\propto g(y_t|x_t)p(x_t|y_{1:t-1}) \tag{13}$$
$$\propto \int g(y_t|x_t)f(x_t|x_{t-1})p(x_{1:t-1}|y_{1:t-1})dx_{1:t-1} \tag{14}$$

Assuming that a particle approximation to $p(x_{1:t-1}|y_{1:t-1})$ has already been generated,

$$p(x_{1:t-1}|y_{1:t-1}) \approx \sum_{i=1}^{N} \delta(x_{1:t-1} - x_{1:t-1}^{(i)}) \tag{15}$$

Then, assuming that $f(x_t|x_{t-1})$ and $g(y_t|x_t)$ can be evaluated pointwise, we generate, for each state trajectory $x_{1:t-1}^{(i)}$, a random sample from a proposal distribution $q(x_t|x_{1:t-1}^{(i)}, y_{1:t})$. Then, the weights w_t of the filtering distribution (eq. 8) can be approximated by

$$w_t^{(i)} \approx \frac{g(y_t|x_t^{(i)})f(x_t^{(i)}|x_{t-1}^{(i)})}{q(x_t^{(i)}|x_{1:t-1}^{(i)}, y_{1:t})} \tag{16}$$

Finally, we perform a multinomial resampling step, such that the probability that $x_t^{(i)}$ is selected is proportional to $w_t^{(i)}$, to obtain an unweighted approximate random draw from the filtering distribution $p(x_t|y_{1:t})$. It is noteworthy that if the resampling step is forgotten, a degeneracy phenomenon can occur. Indeed, after a few iterations, all but one particle will have negligible weight. (Doucet et al., 2000) has shown that the variance of the importance weights can only increase over time, and thus, it is impossible to avoid the degeneracy phenomenon. This degeneracy implies that a large computational effort is devoted to updating particles whose contribution is almost zero. As a result, a resampling step is needed to eliminate particles with small weights and generate a new set $\{x_t^{(i)}\}_i$, which is an i.i.d. (independent and identically distributed) sample from the approximate density $p(x_t|y_{1:t})$, with a resetting of the weights $\{w_t^{(i)}\}_i$ to $1/N$.

1.3.3 Smoothing

The entire smoothing density $p(x_{1:T}|y_{1:T})$ can be factorized as :

$$p(x_{1:T}|y_{1:T}) = p(x_T|y_{1:T}) \prod_{t=1}^{T-1} p(x_t|x_{t+1:T}, y_{1:T}) \tag{17}$$

Using the filter approximation (eq. 8) to $p(x_t|y_{1:t})$ and the Markovian assumptions of the model, we can write,

$$p(x_t|x_{t+1:T}, y_{1:T}) \propto p(x_t|y_{1:t}) f(x_{t+1}|x_t) \approx \sum_{i=1}^{N} w_{t|t+1}^{(i)} \delta(x_t - x_t^{(i)}) \tag{18}$$

with the modified weights

$$w_{t|t+1}^{(i)} = \frac{w_t^{(i)} f(x_{t+1}|x_t^{(i)})}{\sum_{j=1}^{N} w_t^{(j)} f(x_{t+1}|x_t^{(j)})} \tag{19}$$

This revised particle distribution can be used to generate states successively in the reverse-time direction, conditioning upon future states.

1.4 Our contributing work

The main objective of this paper is to propose an alternative formulation of current PLCA models applied to audio signals, replacing the EM algorithm by a more generic parameter estimation algorithm based on a PF method. We call this new algorithm PLCA-PF in the following. The main advantage expected from this new algorithm is to be able to scan the whole parameter space so as to take into account any features of the parameters, and thus overcoming the limitations underlined in our introduction specific to current PLCA models. In regards to prior integration particularly, this new framework allows releasing the constraints on prior mathematical forms and number. This paves the way towards more complete modelings of the multi-faceted information carried by musical signals, covering both time (e.g. tempo and rhythm) and frequency (e.g. note spectra and chords) domains, and the different prior knowledge classes related to musicology, timbre and playing style.

2 Particle Filtering for PLCA

2.1 State space representation

Considering the equations 1 - 2, the PLCA can be expressed as :

$$\begin{aligned} P(x,t) &= \sum_{z} P(z,t)P(x|z) \\ &= \sum_{z_1,\ldots,z_K} P(z_1,\ldots,z_K,t)\prod_{j=1}^{J} P(x_j|z_1,\ldots,z_K) \\ &= \sum_{z_1,\ldots,z_K} P(z_K,t)\prod_{k=1}^{K-1} P(z_k|z_{k+1},\ldots,z_K,t) \\ &\qquad \prod_{j=1}^{J} P(x_j|z_1,\ldots,z_K) \end{aligned} \tag{20}$$

with :

- $z \in Z_1 \times \ldots \times Z_K$ is a vector of K latent components $(z_1,\ldots,z_K)$ associated to a finite subset $Z_k = \{1,\ldots,L_k\}$
- $t \in \{0,\ldots,T\}$ is the time variable
- $x \in X_1 \times \ldots \times X_J$ is a vector of J features $(x_1,\ldots,x_J)$ where $X_j = \{1,\ldots,F_j\}$

In this decomposition, $P(z_K,t)$ can be seen as the activation distribution of the latent variable z_K, $P(z_k|z_{k+1},\ldots,z_K,t)$ as the weight of the variable z_k conditionally to $(z_{k+1},\ldots,z_K)$ and $P(x_j|z_1,\ldots,z_K)$ as the J features basis.

To estimate the set of parameters $p_t = \{P(z_K,t), P(z_k|z_{k+1},\ldots,z_K,t)\forall k \in \{1,\ldots,K-1\}\}$ at each time $t \in \{0,\ldots,T\}$, the model can be rearranged as a state space process

$$p_t \sim f(p_t|p_{t-1}) \tag{21}$$

$$y_t \sim g(y_t|p_t) \tag{22}$$

where f is the transition state density function for p_t defined above and g the observation function of y_t.

2.2 Transition and observation densities

2.2.1 Transition density

Assuming that each latent variable $z_k \in Z_k$ is i.i.d, each marginal vector $P(z_K, t)$ and $P(z_k|z_{k+1}, \ldots, z_K, t)$ can be independently estimated. Recalling that at a given time t, $P(z_K, t)$ and $P(z_k|z_{k+1}, \ldots, z_K, t)$ represent distributions, Dirichlet priors are injected to ensure that their elements belong to $[0, 1]$, as follows, $\forall(z_2, \ldots, z_K) \in Z_2 \times \ldots \times Z_K, \forall k \in \{1, \ldots, K-1\}$

$$P(z_K, t) \sim Dir(\theta_t^1, \ldots, \theta_t^{L_K}) \tag{23}$$

$$P(z_k|z_{k+1}, \ldots, z_K, t) \sim Dir(\delta_k(z_k, \ldots, z_{K-1})_t^1, \ldots, \delta_k(z_k, \ldots, z_{K-1})_t^{L_K}) \tag{24}$$

where θ and δ_k are random variables representing the weight of each component of z_K in $P(z_K, t)$ and $P(z_k|z_{k+1}, \ldots, z_K, t)$. That injection leads to the following hierarchical model

$$\begin{array}{ccc} H_t & \rightarrow & H_{t+1} \\ \downarrow & & \downarrow \\ P_t & & P_{t+1} \\ \downarrow & & \downarrow \\ Y_t & & Y_{t+1} \end{array} \tag{25}$$

with $H_t = (\Theta_t, \Delta_t)$ the new states defined by

$$\Theta_t = \{\theta_t^{z_K}, \forall z_K \in Z_K\} \tag{26}$$

$$\Delta_t = \{\delta_k(z_k, \ldots, z_{K-1})_t^{z_K}, \forall z_K \in Z_K\} \tag{27}$$

where we have defined

$$\theta_{t+1}^{z_K} = \theta_t^{z_K} \times \alpha_t^{z_K}, \alpha_t^{z_K} \sim \phi \tag{28}$$

$$\delta_k(z_k, \ldots, z_{K-1})_{t+1}^{z_K} = \delta_k(z_k, \ldots, z_{K-1})_t^{z_K} \times \gamma_t^{z_K}, \gamma_t^{z_K} \sim \psi_k \tag{29}$$

with ϕ and ψ_k are positive distributions.

2.2.2 Observation density

y_t has been defined as a representation of x at time t. In that state space approach, each component of y_t is represented by the sum of the PLCA model and a white noise, $\forall x \in X_1 \times \ldots \times X_J$,

$$y_t(x) = P(x,t) + V_t = \sum_{z_1,\ldots,z_K} P_t(z_1,\ldots,z_K) \prod_{j=1}^{J} P_t(x_j|z_1,\ldots,z_K) + V_t \quad (30)$$

where $V_t \sim N(0,\sigma^2)$. Denoting $\hat{y}_t$ the vector of components $P(x,t)$, the observation density g follows a normal distribution, i.e. $g \sim N(\hat{y}_t,\sigma^2)$.

2.3 Prior injection

To overcome limits of current music information retrieval systems, a practical engineering solution was to use computational techniques from statistics and digital signal processing allowing the insertion of prior knowledge from different scientific disciplines (e.g. cognitive science, neuroscience, musicology, musical acoustics) (Engelmore and Morgan, 1988; Nawab and Lesser, 1992; Carver and Lesser, 1992; Ellis, 1996). Such systems perform a process of reconciliation between the observed acoustic features and the predictions of an internal model of the data-producing entities in the environment. This approach is close to human experience, who perceive the sound in view of different hierarchical levels of prior knowledge, using a collection of global properties, such as musical genre, tempo, and orchestration, as well as more instrument specific properties, such as timbre. It has been widely applied in many music information retrieval tasks, such as genre recognition and automatic music transcription (Ellis, 1996; Godsmark and Brown, 1999; Bello, 2000; Ryynanen, 2004; Klapuri, 2004; Benetos et al., 2013b).

In mathematical terms, priors are used to sharpen up estimation of model parameters by emphasizing the most likely values in their distributions. This prior integration is performed during state generation, by reweighting each particle value with a corresponding prior gain. Adding prior knowledge on parameters p_t leads up to sample from the posterior distribution,

$$P(p_t|y_t) \sim P(y_t|p_t)P(p_t) \tag{31}$$

where $P(y_t|p_t)$ identifies to the observation density g and $P(p_t)$ the prior knowledge. When the prior and the likelihood are conjugate, sampling from the posterior distribution is rather straightforward. When the posterior does not have a well known form, as it is the case in most real-life applications, computational statistics methods can be introduced to sample from the posterior. The Metropolis-Hasting algorithm (Roberts et al., 1997; Newman and Barkema, 1999; Robert and Casella, 1999), based on Monte Carlo methods, brings a powerful framework to tackle that issue. This algorithm is a random walk that uses an acceptance/rejection rule to converge to the specified target distribution, and proceeds as described in the algorithm 1.

Algorithm 1 Metropolis-Hasting algorithm for prior integration.

1: Draw a starting point p_t^0, for which $P(p_t^0|y_t) > 0$, from a starting distribution $p_0(p_t)$;
2: **for** $q = 1, 2, \ldots$ **do**
3: Sample a proposal p_t^* from a jumping distribution at iteration q, $J_q(p_t^*|p_t^{q-1})$;
4: Calculate the ratio of densities,

$$r = \frac{P(p_t^*|y_t)/J_q(p_t^*|p_t^{q-1})}{P(p_t^{q-1}|y_t)/J_q(p_t^{q-1}|p_t^*)} \tag{32}$$

5: Set

$$p_t^q = \begin{cases} p_t^* & \text{with probability } \min(r,1) \\ p_t^{q-1} & \text{otherwise} \end{cases} \tag{33}$$

6: **end for**

Considering the filtering particle framework defined above, few remarks about the different steps can be highlighted. First, the initial draw is replaced by the PF draw we want to interfere in. Even if the posterior distribution is unknown, the ratio r can be computed as the ratio of the product of the likelihood and the prior since the normalization constant is removed in the ratio.

$$r = \frac{g(y_t|p_t^*)p(p_t^*)/J_q(p_t^*|p_t^{q-1})}{g(y_t|p_t^{q-1})p(p_t^{q-1})/J_q(p_t^{q-1}|p_t^*)} \tag{34}$$

The jumping distribution J_q is chosen as a normal distribution to simplify the ratio computation. Indeed, the symmetry property of the normal distribution involves that J_q can be removed in eq. 34, to get

$$r = \frac{g(y_t|p_t^*)p(p_t^*)}{g(y_t|p_t^{q-1})p(p_t^{q-1}))} \tag{35}$$

3 Application to AMT

We now propose an application of our PLCA-PF framework to the task of Automatic Music Transcription (AMT), and present evaluation results on this task with quantitative comparisons with other state-of-the-art methods.

3.1 Background on AMT

Work on AMT dates back more than 30 years, and has known numerous applications in the fields of music information retrieval, interactive computer systems, and automated musicological analysis (Klapuri, 2004). Due to the difficulty in producing all the information required for a complete musical score, AMT is commonly defined as the computer-assisted process of analyzing an acoustic musical signal so as to write down the musical parameters of the sounds that occur in it, which are basically the pitch, onset time, and duration of each sound to be played. This task of "low-level" transcription, to which we will restrict ourselves in this study, has interested more and more researchers from different fields (e.g. library science, musicology, machine learning, cognition), and been a very competitive task in the Music Information Retrieval community (, 2007) since 2000. Despite this large enthusiasm for AMT challenges, and several audio-to-MIDI converters available commercially, perfect polyphonic AMT systems are out of reach of today's technology (Klapuri, 2004; Benetos et al., 2013b). To overcome these limitations, a practical engineering solution was to use computational techniques from statistics and digital signal processing allowing the insertion of prior knowledge from cognitive science, musicology and musical acoustics (Engelmore and Morgan, 1988; Ellis, 1996). This approach is close to human experience, in which the

perception of sounds is embedded with prior knowledge, using a collection of global properties such as musical genre, tempo, and orchestration, as well as more specific properties, such as the timbre of a particular instrument.

3.2 Acoustic modeling

3.2.1 PLCA formalization

In the audio framework, PLCA views the input magnitude spectrogram of a sound source as a histogram of "sound quanta" across time and frequency, and modeling it as as a linear combination of spectral vectors from a dictionary (Smaragdis et al., 2006). PLCA method is then based on the assumption that a suitably normalized magnitude spectrogram, V, can be modeled as a joint distribution over time and frequency, $P(f,t)$, with f is the log-frequency index and t the time index. This quantity can be factored into a frame probability $P(t)$, which can be computed directly from the observed data (i.e. energy spectrogram), and a conditional distribution over frequency bins $P(f|t)$, as follows

$$P(f,t) = P(t)P(f|t) \tag{36}$$

Spectrogram frames are then treated as repeated draws from an underlying random process characterized by $P(f|t)$. We can model this distribution with a mixture of latent factors related to polyphonic music transcription of single instruments as follows:

$$P(f|t) = \sum_{i,m} P(i|t)P(m|i,t)P(f|i,m) \tag{37}$$

where $P(f|i,m)$ are the spectral templates for pitch $i \in \mathbf{I}$ (with $\mathbf{I}$ the set of pitches, and N_I the number of pitches) and playing mode $m \in \mathbf{M}$ (with $\mathbf{M}$ the set of playing modes, and N_m the number of modes), $P(m|i,t)$ is the playing mode activation, and $P(i|t)$ is the pitch activation (i.e. the transcription). In this paper, the playing mode m will refer to different dynamics of instrument playing (i.e. note loudness). (Smaragdis et al., 2008) extended the PLCA model of eq. 37 by exploiting the fact that in a CQT, a change of fundamental frequency is reflected by a simple frequency translation of its partials, resulting in a shift invariance over log-frequency. The model proposed, called Shift-Invariant PLCA (SIPLCA), then consists in shifting the templates $P(f|i,m)$ over the log-frequency range of the CQT,

thus performing a multi-pitch detection with a frequency resolution higher than MIDI scale. Eq. 37 is re-written as follows

$$P(f|t) = \sum_{i,\delta_f,m} P(i|t)P(m|i,t)P(\delta_f|i,t)P(f-\delta_f|i,m) \tag{38}$$

where δ_f is the pitch shifting factor. To constrain δ_f so that each sound state template is associated with a single pitch, the shifting occurs in a semitone range around the ideal position of each pitch. Thus because we are using in this paper a log-frequency representation with a spectral resolution of 60 bins/octave, i.e. a 20 cent resolution, we have $\delta_f \in$ [-2:2], with N_{Δ_f} the length of this set of values.

$$P(f,t) = \sum_{i,m,\delta_f} A_t(i)B_t(i,m)C_t(i,m,\delta_f)P(f-\delta_f|i,m) \tag{39}$$

In eq. 39, we also identify the different PF arguments, where at a given time t, A_t is a vector of length N_I representing the elements of the pitch activity matrix $P(i,t)$ (equal to $P(t)P(i|t)$, through the Baye's rule), $B_t(i,s)$ is the $N_I \times N_m$ matrix whose coefficients are the weights $P(s|i,t)$, and C_t is the $N_I \times N_m \times N_{\Delta_f}$ tensor corresponding to the spectral weights coefficients $P(\delta_f|i,t,m)$. The spectral shifted templates $P(f-\delta_f|i,m)$ are extracted from isolated note samples using a one component PLCA, and are not updated.

Eventually, as in most spectrogram factorization-based transcription or pitch tracking methods (Grindlay and Ellis, 2011; Mysore and Smaragdis, 2009; Dessein et al., 2010), we use a simple threshold-based detection of the note activations from the pitch activity matrix $P(i,t)$, followed by a minimum duration pruning. The threshold for minimum duration for pruning was set to 50 ms. The use of this simple thresholding method should allow one to better highlight the intrinsic differences from the different AMT systems we will compare.

3.2.2 Particle filter argument

Since, the spectral shifted templates $P(f-\delta_f|i,m)$ are learned, the set of unknown parameters is $\{A_t, B_t, C_t\}$, and $y_t = V_{ft}(.,t) \in [0,1]^F$ denotes the observations. As described in 2.2.1, each marginal vector $A_t(s)$, $B_t(s)$ and $C_t(s,z)$ are independently forecasted through a Dirichlet distribution, as follows

$$A_t \sim Dir(\theta_t^1, \ldots, \theta_t^I) \tag{40}$$

$$B_t(s) \sim Dir(\delta_1(s)_t^1, \ldots, \delta_1(s)_t^I) \tag{41}$$

$$C_t(s, \delta_f) \sim Dir(\delta_2(s, \delta_f)_t^1, \ldots, \delta_2(s, \delta_f)_t^I) \tag{42}$$

Concerning the ϕ, ψ_1 and ψ_2 distributions producing the states $(\theta, \delta_1, \delta_2)$, we opt for a Gamma distribution to obtain a non biased transition and to control the variance of the transition. That choice leads to the following transition rules where $\forall i, s, \delta_f$,

$$\theta_{t+1}^i = \theta_t^i \times \alpha_t^i, \alpha_t^i \sim \Gamma(a^i, b^i) \tag{43}$$

$$\delta_1(s)_{t+1}^i = \delta_1(s)_t^i \times \gamma_t^i, \gamma_t^i \sim \Gamma(c_s^i, d_s^i) \tag{44}$$

$$\delta_2(s, \delta_f)_{t+1}^i = \delta_2(s, \delta_f)_t^i \times \lambda_t^i, \lambda_t^i \sim \Gamma(e_{s,\delta_f}^i, f_{s,\delta_f}^i) \tag{45}$$

with hyperparameters a^i, b^i, c_s^i, d_s^i, e_{s,δ_f}^i and f_{s,δ_f}^i. Conditionally to θ_t and β_t

$$\theta_{t+1}^i \sim \Gamma(a_s^i, \frac{b_s^i}{\theta_t^i}) \tag{46}$$

$$\delta_1(s)_{t+1}^i \sim \Gamma(c_s^i, \frac{d_s^i}{\delta_1(s)_t^i}) \tag{47}$$

$$\delta_1(s, \delta_f)_{t+1}^i \sim \Gamma(e_{s,\delta_f}^i, \frac{f_{s,\delta_f}^i}{\delta_2(s, \delta_f)_t^i}) \tag{48}$$

Recalling that the mean of θ_{t+1}^i is $\frac{a_s^i \theta_t^i}{b_s^i}$, a non biased transition for θ involves that $a_s^i = b_s^i$ since $E(\theta_{t+1}^i) = \theta_{t+1}^i$ is expected. Under the same argument, $c_s^i = d_s^i$ and $e_{s,\delta_f}^i = f_{s,\delta_f}^i$. Figure 1 provides an example of piano-roll transcription output obtained with our PF-PLCA system.

3.3 Prior knowledge integration

A musical signal is highly structured, in both time and frequency domains. In time domain, tempo and beat specify the range of likely note transition times. In the frequency domain, as audio signals are both additive and oscillatory (musical objects in polyphonic music superimpose and not conceal each other), several notes played at the same time form chords, or polyphony[2],

[2]Here polyphonic music refers to a signal where several sounds occur simultaneously. Whereas in monophonic signals, at most one note is sounding at a time.

merging their respective spectral structures. When designing priors for an AMT system, one basically aims to help the system figuring out "which notes are present at time t" and "by which ones they will be followed". These two types of information belong respectively to frame-wise spectral priors (e.g. sparseness, spectrum modeling including inharmonicity (Rigaud et al., 2013)) and to frame-to-frame temporal priors (e.g. harmonic content transitions, smoothing of spectrum envelop), and will both be developed into our PLCA-PF framework.

In transcription systems with a general application (Emiya et al., 2010; Fuentes et al., 2013; Benetos and Dixon, 2013), prior knowledge is generally incorporated with no regard to their musical/physical sense, but with the sole preoccupation of convergence optimization and enhancement of transcription results on a specific musical corpus. As a consequence, priors mostly take the form of a single constant factor, set arbitrarily after simulation experiments. Here, we propose a more complex modeling of musical signal with explicit music-related knowledge.

To do so, we quantify relations of influence between the different pitches of the instrument pitch range, either within a frame (for spectral priors), or from one frame to the next one (for temporal priors). These priors will then take the form of a matrix S of size N_I^2, which quantifies average relations of influence between the $N_I \times N_I$ couples of different pitches, and is defined as follows, $\forall (i,j) \in \{1, \cdots, N_I\}^2$,

$$S = \begin{pmatrix} S(1,1) & \cdots & S(1,N_I) \\ \vdots & S(i,j) & \vdots \\ S(N_I,1) & \cdots & S(I,N_I) \end{pmatrix} \tag{49}$$

The PLCA-PF framework allows a general insertion of this matrix through the term $P(p_t)$ of eq. 31, to which we can give the following form

$$P(p_t) \propto \exp(-p_t' S K_{p_t}) \tag{50}$$

where $p_t \in \{A_t, B_t(s), C_t(s,\delta_f)\}$, $\forall (s,\delta_f) \in \{1,\ldots,S\} \times \{1,\ldots,\Delta_f\}$, and K_{p_t} is a vector of length N_I associated to p_t.

It is noteworthy that simpler modeling of prior knowledge, such as a simple pitch-dependent vector, can also take the form of a diagonal matrix S of size N_I^2, with the vector values put into this diagonal (the zero-coefficients of S provide an unitary prior value which does not affect particle weights).

3.3.1 Sparse priors

During the multi-pitch estimation step of an AMT process, a too much large number of non-zero activation scores is often observed, making the operation of "finding the right notes" more difficult. In order to overcome this flaw, a sparseness prior can reduce the number of active notes per frame in selecting the most salient ones. Previous works mostly use pitch-independent sparse prior in PLCA-EM algorithms. (Fuentes et al., 2013) compute a sparseness prior $P(A_t)$ to constrain the impulse distribution A_t, as follows

$$P(A_t) \propto \exp\left(-2\beta\sqrt{J}||A_t)||_{1/2}\right) \tag{51}$$

with $||A_t||_{1/2} = \sum_i \sqrt{A_t(i)}$ and β a positive hyperparameter indicating the strength of the prior. With this prior, a numerical fixed point algorithm is required to obtain a solution with the EM algorithm. Other works (Grindlay and Ellis, 2011; Benetos and Dixon, 2011, 2013) impose sparsity on the pitch activity matrix and the pitch-wise source contribution matrix by modifying EM equations.

The common point to all these EM-based sparse priors is that they are pitch-independent, and rely on hyper-parameters, which are either arbitrary set and/or optimized on a given sound dataset. In this paper, we define sparse priors informed by explicit musical acoustics related knowledge. Musically, the occurrence of simultaneous notes can result either from "acoustic polyphony", or from "musical polyphony". "Acoustics polyphony" is strongly related to the timbre of the instrument, and more precisely to the physical phenomena of mutual resonances and note persistence. Although this type of polyphony is an integral part of instrument timbre, it represents a noise signal added to the actual played note from the point of view of music transcription. For what concerns "musical polyphony", it corresponds to the note combinations played by the musician and intended by a composer with a proper polyphonic writing. It directly provides useful information about which notes are commonly played simultaneously in a musical piece. Prior knowledge can be learned from both of these polyphonic origins for an instrument repertoire, studying respectively the timbre of the instrument or the frame-wise musical characteristics of the repertoire.

A first sparse prior P_{spa1} on note mixture likelihood has then been defined. Following the form of matrix S (eq. 49), each coefficient is computed by a frame-wise counting of the pitches j played simultaneously to pitch i,

from our training MIDI transcripts (see Sec. 3.5.2 for details on the sound database). We propose a second sparse prior P_{spa2} on mutual resonances. For strings on a bowed, plucked, or hammered instruments, mutual resonances result from sympathetic strings, which vibrate (and thereby sound a note) in sympathetic resonance with the note sounded near them by some other agent. Here, to compute each coefficient $S(i,j)$, we used two datasets of isolated notes, a first one composed of free-resonating notes, and a second one in which all strings were muted excepting the played one. For each note sample of these two datasets, the spectrum was computed with a FFT using a 4096-sample Hamming window after the onset, unitary normalized and labelled $X_{(d,i)}$ for pitch i and dataset d (equal to 1 or 2). We then used the algorithm 2 to get the scores S(i,j).

Algorithm 2 Computation of coefficients S(i,j).

for For each pitch i ∈ I **do**
 $\tilde{X}_i = ||X_{(1,i)} - X_{(2,i)}||$
 Binary thresholding of $\tilde{X}_i$, i.e.

$$\tilde{X}_i(f) = \begin{cases} 1 & \text{for } f = arg(\tilde{X}_i \geq 0.5) \\ 0 & \text{otherwise} \end{cases}$$

 for Each pitch j ∈ $I, j \neq i$ **do**
 S(i,j) = $X_{(2,j)} \cdot \tilde{X}_i$, with [·] the element-wise product
 end for
end for

These two priors P_{spa_1} and P_{spa_2} are represented in figure 2 through their respective matrices S_{spa_1} and S_{spa_2}. Eventually, for this sparse type prior, the set of prior parameters p_t in eq. 50 is equal to A_t and K_{p_t} is set to A_t , which becomes

$$P_{spa}(A_t) \propto \exp(-A_t' S_{spa} A_t) \tag{52}$$

For prior P_{spa1}, before injecting it in eq. 52, we normalized its matrix S_{spa1} with the operator Π (eq. 53) defined as

$$\Pi(x) = 1 - \frac{x}{max(x)} \tag{53}$$

as we need knowledge rejecting the hypothesis of certain pitch combinations.

3.3.2 Sequential priors on harmonic transitions

Many previous works (Poliner and Ellis, 2007; Grindlay and Ellis, 2011; Benetos and Dixon, 2013) on AMT have used sequential priors to model each pitch activity/inactivity phases, which is done using two-state on/off HMMs for each of them during a post-processing stage. This operation performs a time filtering of note detection decision, which mainly avoids a lot of single miss errors and smooths note boundaries. But musically, the information is very restricted, as it consists only in knowing how long a given pitch note remains active, which can result from both playing techniques of the musician and vibratory properties of the instrument.

The sequential prior we present in this paper is defined as the probability to switch between two successive mixtures of notes in a musical piece. These transition probabilities are determined by sampling the training MIDI transcripts at the precise times corresponding to the analysis frames of the activation matrix, and just checking for the presence of a note in each frame. These probabilities give us a global view of the usual and unusual harmonic transitions of an instrument repertoire. Original mixtures, i.e. those not encountered during the training phase, get a likelihood weighted accordingly to the Witten-Bell discounting algorithm (Witten and Bell, 1991). This prior P_{tra} is represented in figure 2 through their matrix S_{tra}. Eventually, for this sequential type prior, p_t in eq. 50 is equal to A_t and K_{p_t} is set to A_{t-1}, which leads to

$$P_{tra}(A_t) \propto \exp(-A_t' S_{tra} A_{t-1}) \tag{54}$$

3.3.3 Prior combination

The PLCA-PF framework offers an easy-to-implement unifying way of integrating priors from both time and frequency domains. In this framework, priors are injected during the filtering process through eq. 31, and modify the parameters without disturbing their generation. In the set of parameters p_t, the independence between each parameter p_t^n leads to

$$P(p_t) \propto \prod_n P(p_t^n) \tag{55}$$

Within a defined parameter p_t^n, the general prior $P(p_t^n)$ can be seen as the product of the different priors P_{prior}^n associated to p_t^n

$$P(p_t^n) \propto \prod_{prior} P_{prior}^n(p_t^n) \tag{56}$$

Using equations 50, 55 and 56, we combine the different priors, characterized by their respective matrices S_{prior}, as follows

$$P(p_t) \propto \exp(-\sum_n \sum_{prior} p_t^{n\prime} S_{prior}^n K_{prior}^n) \tag{57}$$

3.4 Practical implementation

As a time-frequency representation, all input signals sampled undergo a Q-constant with 60 bins/octave, with window size of 23 ms (1024 coefficients at 44.1-kHz sampling rate) and a 50 % hop, which is adequate for the tonal part of the signal. We now present the different algorithms implemented in our PF framework.

3.4.1 Filtering algorithm

Such as in (Fong et al., 2002), we develop a generic PF algorithm assuming that the proposal distribution $q(x_t|x_{1:t-1}^{(i)}, y_{1:t}) = f(x_t|x_{t-1}^{(i)})$, as detailed in the pseudo-algorithm 3. The resample step is processed as detailed in the pseudo-algorithm 4.

Algorithm 3 Generic particle filtering with multinomial resampling

1: Let $f(x_1|x_0) = f(x_1)$ be state prior distribution. Then for $t = 1$ to T :

2: $\forall i \in \{1, \cdots, N\}$, generate N samples from the proposal $q(x_t|x_{1:t-1}^{(i)}, y_{1:t}) = f(x_t|x_{t-1}^{(i)})$,

$$x_t^{(i)} \sim f(x_t|x_{t-1}^{(i)})$$

3: $\forall i \in \{1, \cdots, N\}$, evaluate the importance weights and normalise

$$w_t^{(i)} \propto g(y_t|x_t^{(i)}), \sum_{i=1}^{N} w_t^{(i)} = 1$$

4: Resample $\{x_t^{(i)}; i = 1, \cdots, N\}$ N times with replacement.

Algorithm 4 Multinomial resampling

1: Initialize the CDF : $c_1 = 0$
2: $\forall i \in \{2, \cdots, N\}$
 - Construct CDF : $c_i = c_{i-1} + w_t^{(i)}$
3: Start at the bottom of the CDF : $i = 1$
4: Draw a starting point : $u_1 \sim U[0, \frac{1}{N}]$
5: $\forall j \in \{1, \cdots, N\}$
 - Move along the CDF : $u_j = u_1 + \frac{j-1}{N}$
 - While $u_j > c_i$, $i = i + 1$
 - Assign sample : $x_t^{(j)*} = x_t^{(i)}$
 - Assign weight : $w_t^{(j)} = 1/N$
6: Return $\{x_t^{(k)*}, w_t^k\}_{k=1}^N$.

3.4.2 Smoothing algorithm

After having generated weighted particles $\{x_t^{(i)}, w_t^{(i)}; i = 1, \cdots, N, t = 1, \cdots, T\}$ from the PF, the smoothing algorithm proceeds as detailed in the pseudo-algorithm 5.

Algorithm 5 Generic particle smoother

1: Choose $\widetilde{x}_T = x_T^{(i)}$ with probability $w_T^{(i)}$.
2: For $t = T - 1$ to 1 :
 - Calculate $w_{t|t+1}^{(i)} \propto w_t^{(i)} f(\widetilde{x}_{t+1}|x_t^{(i)})$ for each $i = 1, \cdots, N$;
 - Choose $\widetilde{x}_t = x_t^{(i)}$ with probability $w_{t|t+1}^{(i)}$.
3: $\widetilde{x}_{1:T} = (\widetilde{x}_1, \widetilde{x}_2, \cdots, \widetilde{x}_T)$ is an approximate realisation from $p(x_{1:T}|y_{1:T})$.

3.5 Evaluation procedure

3.5.1 Evaluation AMT systems

To evaluate comparatively transcription performance of our PLCA-PF algorithm, we tested different algorithms on the same test datasound. Table 1 provides an overview of these algorithms. HALCA is short for the

Harmonic Adaptive Latent Component Analysis algorithm[3] (Fuentes et al., 2013). Here, each note in a constant-Q transform is locally modeled as a weighted sum of fixed narrowband harmonic spectra, spectrally convolved with some impulse that defines the pitch. All parameters are estimated by means of the EM algorithm, in the PLCA framework. This algorithm was recently evaluated by MIREX and obtained the 2^{nd} best score (, 2007, 2^{nd} best scores in the Multi-Pitch Estimation task, 2009-2012). The algorithm PLCA-EM is the algorithm proposed by (Benetos et al., 2013a), whose main characteristics is its use of pre-defined templates, allowing them to avoid updating the templates in the maximization step of the EM algorithm. This algorithm is also state-of-the-art (, 2007, 1^{st} best scores in the Multi-Pitch Estimation task, 2009-2012). PLCA-DAEM is the same as PLCA-EM, only replacing the EM algorithm by a DAEM algorithm. (Cheng et al., 2013) observed significant improvements on their transcriptions through this modification.

Method name	References	Parameter estimation algorithm	Priors
HALCA	(Fuentes et al., 2013)	EM	Sparsity + Continuity + Unimodal
PLCA-EM	(Benetos et al., 2013a)	EM	Sparsity
PLCA-DAEM	(Cheng et al., 2013)	DAEM	Sparsity
PLCA-PF	Proposed	PF	
PLCA-PF + priors	Proposed	PF	P_{spa1}, P_{spa2}, P_{tra}

Table 1 – Recapitulative table of the AMT systems tested in our simulation experiments.

3.5.2 Musical corpus

To test our AMT system and train the sparse priors proposed, we need three different sound corpus: audio musical pieces of an instrument repertoire, the corresponding scores in the form of MIDI files, and a complete dataset of isolated notes for this instrument. We will use two different decay instruments for evaluation, namely the classical piano and the *marovany* zither from Madagascar. For piano, audio data was extracted from the MAPS database (Emiya et al., 2010), which is composed of high-quality note samples and recordings from a real upright piano, whose MIDI scores have been automatically compiled using the Disklavier technology. For the marovany instrument, sound templates and musical pieces were extracted from personal

[3]Codes are available at `http://www.benoit-fuentes.fr/`.

recordings made in our laboratory. Pieces were transcribed with an original multi-sensor retrieval system (Cazau et al., 2013).

From these sound databases, we extracted different sets of training and test data, as our prior must be trained using automatically generated knowledge from MIDI files and template datasets. To do so, we first divided each musical pieces into 15-second sequences, which provided us with a total of 1.2 and 0.83 hours of audio, respectively for the piano and marovany datasets. Within each dataset, the musical sequences were randomly split into training and testing sequences, using by default 30 % of sequences for testing, and the 70% remaining ones for training. In our simulation experiments, this procedure is repeated five times, and an average is computed on the resulting scores. To prevent any overfitting of our data, we carefully distinguished between training and test data. Especially, sound templates used in the PLCA model were extracted from an instrument model different from the one used in recordings. Also, sequences used to train the priors were not used for evaluation.

3.5.3 Error metrics

For assessing the performance of our proposed transcription system, we adopt a note-oriented approach, according to which a note event is assumed to be correct if it fills the condition that its onset is within a 50 ms range from a ground-truth onset (i.e. the standard tolerance commonly used (Bello et al., 2005; , 2007)). Such a tolerance level is considered to be a fair margin for an accurate transcription, although it is far more tolerant than human ears would, as we remind that those are able to distinguish between two onsets as close as 10 ms apart (Moore, 1997). Evaluation metrics are defined by equations 58-60 (, 2007), resulting in the note-based recall (TPR), precision (PPV) and F-measure (the harmonic mean of precision and recall) :

$$\text{TPR} = \frac{\sum_{n=1}^{N} \text{TP}[n]}{\sum_{n=1}^{N} \text{TP}[n] + \text{FN}[n]} \tag{58}$$

$$\text{PPV} = \frac{\sum_{n=1}^{N} \text{TP}[n]}{\sum_{n=1}^{N} \text{TP}[n] + \text{FP}[n]} \tag{59}$$

$$\text{F-measure} = \frac{2.\text{PPV}.\text{TPR}}{\text{PPV} + \text{TPR}} \tag{60}$$

where N is the total number of notes, and TP, FP and FN scores stand for the well-known True Positive, False Positive and False Negative detections. The recall is the ratio between the number of relevant and original items; the precision is the ratio between the number of relevant and detected items; and the F-measure is the harmonic mean between precision and recall. For all these evaluation metrics, a value of 1 represents a perfect match between the estimated transcription and the reference one.

3.6 Results and discussion

We present in the following simulation experiments on parameter initialization dependency and transcription performance, comparing our proposed PF-based algorithm with three other state-of-the-art algorithms (see table 1).

3.6.1 Computational time

Figure 3 shows the computational time of our PLCA-PF system on a 15-s test musical sequence against the number of particles. Both the computational time and transcription performance increase with the number of particles used. It is well-known that filtering particle is very demanding in computation time, which increases exponentially with the number of particles. In its current form, our system is not very efficient computationally, as it produces a transcription in about 50 at real time on a PC computer (e.g., for a 15 sec recording it requires 12.5 mins) with 2000 particles. Increasing the number of particles also increases transcription accuracy, rising the average F-measure with gains as high as 14 % between 10 and 1000 particles, which begins to stagnate after 5000 particles. This tendency remains observed regardless the instrument repertoire. A trade-off between computation time and transcription precision must then be considered.

3.6.2 Parameter initialization dependency

Figure 4 compares dependency of transcription outputs on parameter initialization. These results are computed from 40 simulation trials on one test sequence with a randomized parameter initialization. Black bars indicate the reference scores of each system, obtained with an uniform initialization of parameters, as it is commonly done by default (e.g. (Fuentes et al., 2013,

Paragraph V.A.1)). We observe that the proposed systems including filtering particle are globally more robust to parameter initialization, in comparison to other EM or DAEM -based systems, which present an important variability in the average F-measure (e.g., ±2.7% the PLCA-EM model for the *marovany* repertoire. Then, depending on the sound dataset under evaluation, the set of initialization parameters may be sub-optimal, with performance losses rising as high as 5 % in the average F-measure. In definitive, making AMT systems less dependent on data should favour their generalization to the diversity of music.

3.6.3 Transcription performance

Table 2 compares transcription performance of our different AMT systems through the different error metrics, for the piano and the *marovany* repertoires, respectively. Both of them present a complex polyphony structure. The *marovany* repertoire is characterized by fast arpeggios, with an ample halo-like sound with rich overtones due to the complex resonating behavior of the instrument. The classical piano repertoire presents more complex and richer chord transitions, with different playing techniques and dynamics which interfere continuously on the timbre of the instrument. For both of these repertoires, the improvements brought by PF-based systems in transcription performance are likely related to the EM limitations evoked in our Introduction, which make EM-based algorithms less efficient in finding active notes in complex polyphonic signals due to problems of local maxima convergences. Also, information from standard deviations shows that our proposed algorithm presents the minimum value for the standard deviation (2.3 %), which implies a higher robustness to transcribe polyphonic signals with different musical features.

	Piano			Marovany		
Methods	TPR	PPV	F-measure	TPR	PPV	F-measure
HALCA	55.2	59.7	57.3	55.1	57.6	56.3
PLCA-EM	57.2	62.8	59.8	55.6	59.7	57.6
PLCA-DAEM	59.1	62.9	60.9	56.1	58.3	57.2
PLCA-PF	57.9	62.2	59.9	57.7	60.1	58.8
PLCA-PF + priors	61.2	62.5	61.8	58.2	60.9	59.5

Table 2 – Mean transcription error metrics (in %) for the piano recordings with our different AMT systems.

4 Conclusion

Current PLCA-based systems for AMT use the well-known EM algorithm to estimate the model parameters. This algorithm presents well-known inherent defaults (local convergence, initialization dependency), making EM-based systems limited in their applications to AMT, particularly in regards to the mathematical form and number of priors. To overcome such limits, we have developed in this paper a different estimation framework based on Particle Filtering, which consists in sampling the posterior distribution over larger parameter ranges. This framework proves to be more robust in parameter estimation, more flexible and unifying in the integration of prior knowledge in the system. It provides the abilities of injecting more complex musicological knowledge, as well as combining simultaneously a theoretically infinite number of priors. Our proposed Particle-Filtering systems achieve promising rankings in terms of accuracy rate, and further experimentations will be necessary to confirm these preliminary results.

Acknowledgements

At the risk of omitting some relevant names, the authors would like to especially thank March Chemillier (CAMS-EHESS) for his help in recording the *marovany*, and Laurent Quartier (LAM-UPMC) for technical supports.

References

(2007), M. (**2011**). "Music information retrieval evaluation exchange (mirex)." Available at http://music-ir.org/mirexwiki/ (date last viewed January 9, 2015).

Andrieu, C., Davy, M., and Doucet, A. (**2003**). "Efficient particle filtering for jump markov systems. application to time-varying autoregressions." IEEE Trans. on Signal Proc., **51**, 1762–1770.

Bello, J.P., Daudet, L., Abdallah, S., Duxbury, C., Davies, M., and Sandler, M.B. (**2005**). "A tutorial on onset detection in music signals." IEEE Trans. on Speech and Audio Proc., **13**, 1035–1047.

Bello, J. P.and Sandler, M.B. (**2000**). "Blackboard systems and topdown processing for the transcription of simple polyphonic music." In *Proceedings of the International Conference on Digital Audio Effects (DAFx)*.

Benetos, E., Cherla, S., and Weyde, T. (**2013**a). "An efficient shift-invariant model for polyphonic music transcription." In *6th Int. Workshop on Machine Learning and Music, Prague, Czech Republic*.

Benetos, E. and Dixon, S. (**2011**). "Multiple-instrument polyphonic music trancription using a convolutive probabilistic model." In *Proc. 8th Sound and Music Computing Conf.* pp. 19–24.

Benetos, E. and Dixon, S. (**2013**). "Multiple-instrument polyphonic music transcription using a temporally constrained shift-invariant model." J. Acoust. Soc. Am., **133**, 1727–1741.

Benetos, E., Dixon, S., Giannoulis, D., Kirchhoff, H., and Klapuri, A. (**2013**b). "Automatic music transcription: Challenges and future directions." J. of Intelligent Information Systems, **41**, 407–434.

Carver, N. and Lesser, V. (**1992**). *Symbolic and Knowledge-Based Signal Processing* (New York: Prentice Hall), chap. Blackboard systems for knowledge-based signal understanding.

Cazau, D., Chemillier, M., and Adam, O. (**2013**). "Information retrieval of marovany zither music with an original optical-based system." In *Proceedings of DAFx 2013, Maynooth, Ireland*. pp. 1–6.

Cheng, T., Dixon, S., and Mauch, M. (**2013**). "A deterministic annealing algorithm for automatic music transcription." In *14th International Society for Music Information Retrieval Conference, Curitiba, PR, Brazil*.

Dempster, A.P., Laird, N.M., and Rubin, D.B. (**1977**). "Maximum likelihood from incomplete data via the em algorithm." Journal of the Royal Statistical Society, Series B, **39**, 1–38.

Dessein, A., Cont, A., and Lemaitre, G. (**2010**). "Real-time polyphonic music transcription with nonnegative matrix factorization and beta-divergence." In *11th International Society for Music Information Retrieval Conference, Utretcht, Netherlands*. pp. 489–494.

Doucet, A., Godsill, S., and Andrieu, C. (**2000**). "On sequential monte carlo sampling methods for bayesian filtering." Statistics and Computing, **10**, 197–208.

Ellis, D.P.W. (**1996**). "Prediction-driven computational auditory scene analysis." Ph.D. thesis, Massachusetts Institute of Technology.

Emiya, V., Badeau, R., and Richard, G. (**2010**). "Multipitch estimation of piano sounds using a new probabilistic spectral smoothness principle." IEEE Trans. on Audio, Speech, Lang. Proc., **18**, 1643–1654.

Engelmore, R. and Morgan, A., eds. (**1988**). *Blackboard Systems* (Addison-Wesley Longman Publishing (Boston, MA, USA)), 602 pp.

Févotte, C. and Godsill, S. (**2006**a). "A bayesian approach for blind separation of sparse sources." IEEE Trans. Audio Speech Language Processing, **14**, 2174–2188.

Févotte, C. and Godsill, S. (**2006**b). "Sparse linear regression in unions of bases via bayesian variable selection." IEEE Signal Processing Letters, **13**, 441–444.

Févotte, C., Torrésani, B., Daudet, L., and Godsill, S. (**2008**). "Sparse linear regression with structured priors and application to denoising of musical audio." IEEE Trans. Audio Speech Language Processing, **16**, 174–185.

Fong, W., Godsill, S., Doucet, A., and West, M. (**2002**). "Monte carlo smoothing with application to audio signal enhancement." IEEE Transactions on Signal Processing, **50**, 438–449.

Fuentes, B., Badeau, R., and Richard, G. (**2013**). "Harmonic adaptive latent component analysis of audio and application to music transcription." IEEE Trans. on Audio Speech Lang. Processing, **21**, 1854–1866.

Godsmark, D. and Brown, G.J. (**1999**). "A blackboard architecture for computational auditory scene analysis." Speech Communication, **27**, 351–366.

Grindlay, G. and Ellis, D.P.W. (**2010**). "A probabilistic subspace model for multi-instrument polyphonic transcripton." In *11th International Society for Music Information Retrieval Conference, Utrecht, Netherlands*.

Grindlay, G. and Ellis, D.P.W. (**2011**). "Transcribing multi-instrument polyphonic music with hierarchical eigeninstruments." IEEE J. Sel. Topics Signal Proc., **5**, 1159–1169.

Hoffman, M.D., Blei, D.M., and Cook, P.R. (**2009**). "Finding latent sources in recorded music with a shift-invariant hdp." In *12th Int. Conference on Digital Audio Effects (DAFx-09), Como, Italy*.

Hofmann, T. (**1999**). "Probabilistic latent semantic indexing." In *22th Annual International SIGIR Conference on Research and Development in Information Retrieval*.

Klapuri, A. (**2004**). "Automatic music transcription as we know it today." J. of New Music Research, **33**, 269–282.

Moore, B.C.J. (**1997**). *An Introduction to the Psychology of Hearing* (New York: Academic), 441 pp.

Mysore, G.J. and Smaragdis, P. (**2009**). "Relative pitch estimation of multiple instruments." In *International Conference on Acoustical Speech and Signal Processing, Taipei, Taiwan*. pp. 313–316.

Nawab, S.H. and Lesser, V. (**1992**). *Symbolic and Knowledge-Based Signal Processing* (New York: Prentice Hall), chap. Integrated processing and understanding of signals.

Newman, M.E. and Barkema, G.T. (**1999**). *Monte Carlo Methods in Statistical Physics* (USA: Oxford University Press).

Poliner, G. and Ellis, D. (**2007**). "A discriminative model for polyphonic piano transcription." J. on Advances in Signal Proc., **8**, 1–9.

Rigaud, F., Falaize, A., David, B., and Daudet, L. (**2013**). "Does inharmonicity improve an nmf-based piano transcription model ?" In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 11–15.

Robert, C.P. and Casella, G. (**1999**). *Monte Carlo Statistical Methods* (Springer Science+Business Media New York).

Roberts, G.O., Gelman, A., and Gilks, W.R. (**1997**). "Weak convergence and optimal scaling of random walk metropolis algorithms." Ann. Appl. Probab., **7 (1)**, 110–120.

Ryynanen, M. (**2004**). "Probabilistic modelling of note events in the transcription of monophonic melodies." Master's thesis, Tampere University.

Smaragdis, P., Raj, B., and Shanshanka, M. (**2006**). "A probabilistic latent variable model for acoustic modeling." In *Neural Information Proc. Systems Workshop, Whistler, BC, Canada*.

Smaragdis, P., Raj, B., and Shashanka, M. (**2008**). "Sparse and shift-invariant feature extraction from non-negative data." In *International Conference Acoustical Speech and Signal Processing, Las Vegas, NV*. pp. 2069–2072.

Ueda, N. and Nakano, R. (**1998**). "Deterministic annealing em algorithm." Neural Networks, **11**, 271–282.

Vermaak, J., Andrieu, C., and Doucet, A. (**2000**). "Particle filtering for non-stationary speech modelling and enhancement." In *6th International Conference on Spoken Language Processing*. pp. 594–597.

Witten, I.H. and Bell, T.C. (**1991**). "The zero-frequency problem: estimating the probabilities of novel events in adaptive text compression." IEEE Transactions on Information Theory, **37**, 1085–1094.

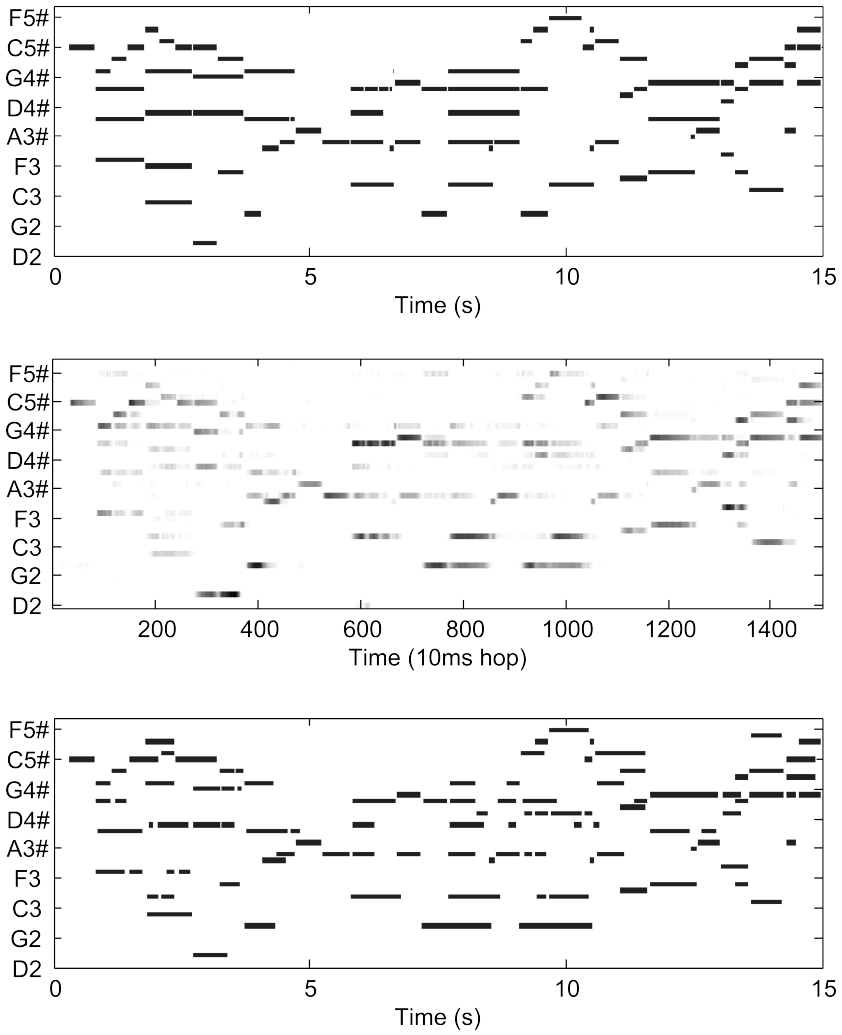


Figure 1 – Illustration of different stages of our PF-PLCA system on a test musical sequence, with from top to bottom: ground truth, pitch activity matrix P(i,t) and piano-roll transcription output.

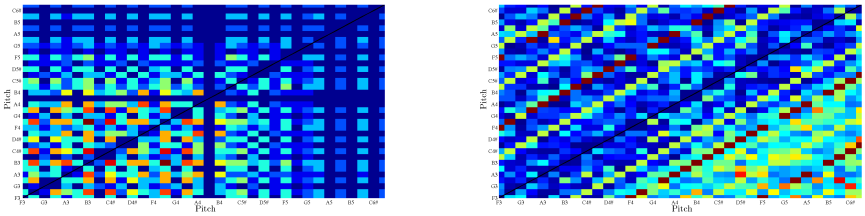


Figure 2 – Illustration of the inter-pitch influence matrix S for chord content (on the left) and sympathetic resonances (on the right).

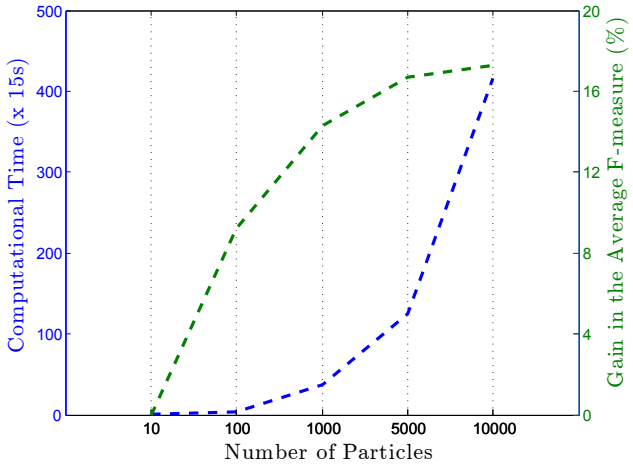


Figure 3 – Computational time of our PLCA-PF system on a 15-s test musical sequence against the number of particles.

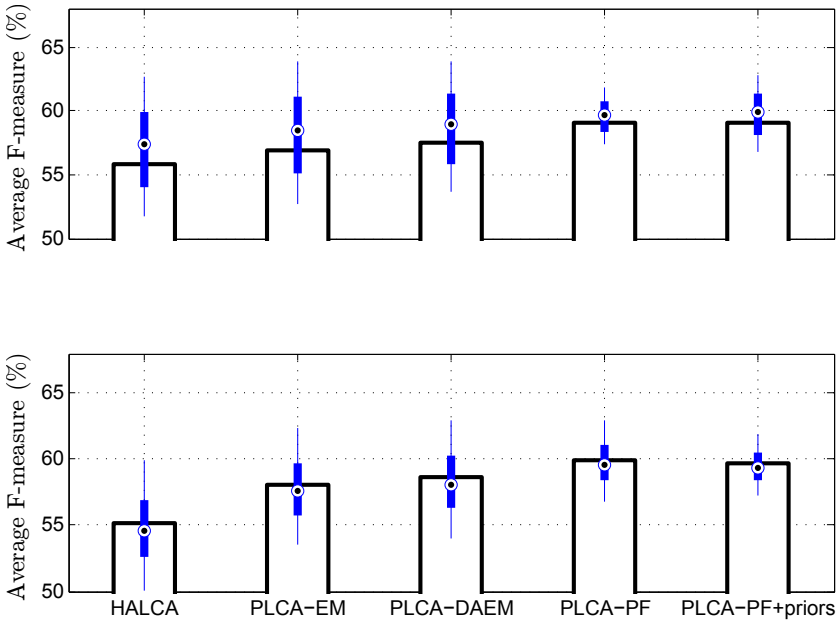


Figure 4 – Variances in transcription performance using random initialization of system parameters, for the piano (top graph) and the marovany (bottom graph) repertoires. Black bars indicate the reference scores of each system, obtained with an uniform initialization of parameters, as it is commonly done by default (e.g. (Fuentes et al., 2013, Paragraph V.A.1)).